_lognostics

Simulating recovery from bilingual aphasia.
Paul Meara

This paper describes some work that I and my colleague Alison Wray have been doing on descriptive models of bilingual lexicons. This work is rather different from the mainstream work that has appeared in this area over the last few years - eg. de Bot's adaptation of Levelt's model, or the work of de Groot and her colleagues (de Bot 1992; Levelt 1989; de Groot 1995). Indeed, our work is in many ways a deliberate reaction against the complexity of these models. Although the approaches we have mentioned differ in their details, they all share a set of common working assumptions. These assumptions are, basically, that "lexicons" can be shown to produce characteristic behaviour patterns, and that these behaviour patterns can be explained by devising formal models of how the lexicon is structured. If the model we develop predicts behaviours which actually occur, then it is taken to be a good model. There is, of course, nothing wrong with this approach. Indeed, some of the more recent models of this type seem to be extraordinarily accurate (eg. Dijkstra and de Smedt 1996). Nonetheless, it is difficult to avoid the feeling that this high level of success has been bought at some cost: not even their strongest supporters would describe current models of the mental lexicon as simple.

This has made us wonder whether it might be interesting to explore the properties of very much simpler models of lexical organisation. Our strategy in this work has been to build non-complex models - models which are not designed to explain particular pre-conceived problems - and to test what properties emerge spontaneously from them. That is, instead of designing models of the lexicon so that they must produce a set of predetermined desired behaviour patterns, we work with very simple models, examine their properties, and see if these properties have any analogues with the behaviour patterns produced by real lexicons.

The models that I am going to describe in this paper are so simple that many readers will be surprised that they can have any bearing at all on the a question as complex as the organisation of real mental lexicons. As we shall see, however, some surprising things can emerge from beginnings that look decidedly unpromising.

Our model lexicons belong to a set of cellular automata known as k=2 random boolean networks (Kauffman 1991; 1993). Stripping away the jargon, this means that our "lexicons" consist of a set of units which it is convenient to think of as words. Each "word" is connected randomly to two other words in the system. Each word can be in one of two states, ON or OFF. We usually think of an ON word as one which is currently active or available for production, while an OFF word is a passive vocabulary item, available for production only if it is turned ON by external stimulation. The current state of each word depends on the state of the two words that it is directly connected to. In the models described here, words turn ON or OFF according to a set of simple Boolean rules. Some words turn ON if both of their connections are ON, while some words turn ON if one or both of their connections are ON. These processes are described in more detail in figure 1. The important points to note are that the interconnections between the words in these models are both sparse and random. Only a bare minimum of structure has been built into the models - just enough to distinguish them from unstructured lists of words. (See Appendix one for a more detailed description of random Boolean networks).

In an earlier paper (Meara 1996) I showed that randomly organised minimal networks of this kind had some odd properties which were reminiscent of some of the behaviours that we find in real lexicons. In particular, the randomised networks seem to provide fairly natural accounts of the way words in real lexicons shift between active and receptive states. There is, for example, no need to set up a separate mechanism to account for the "passive to active shift" in these models; the distinction between active productive and receptive passive vocabulary just emerges as a property of the way these networks function. Similarly, the network model provides a possible explanation for some odd data on lexical attrition (Grendel 1993; Weltens and Grendel 1993): once established, Boolean networks are remarkably stable, and up to a point they can repair themselves.

It is also possible to design "bilingual" networks. In our simulations, a bilingual network is constructed by building two separate, independent networks. We then alter a small fraction of the random connections so that some words have one same-language random connection, and one cross-language connection. Like the monolingual models, these bilingual models also appear to have properties which mirror the behaviours of real bilingual speakers (Meara 1996). Like the monolingual networks described earlier, our bilingual networks normally stabilise into a steady state. Typically, this steady state will favour one language over the other, in the sense that one language tends to have a greater number of ON units than the other one does. It is convenient to think of this initially dominant language as the L1, and the weaker language as the L2. This stable state state can be easily disturbed, however. If the weaker language is stimulated by randomly turning words ON, then the entire network can sometimes move rapidly towards a new equilibrium state. When this happens, large numbers of L1 words become inactive, and large numbers of L2 items become active, so that the L2 system as a whole becomes quickly dominant. This pattern of dominance holds as long as the external stimulation is maintained. Once it stops, the system rapidly returns to its stable resting configuration, in which the L1 dominates. This sort of property looks suspiciously like a mechanism that would allow bilingual lexicons to switch rapidly between the languages that made them up. We were also able to show that a randomised model involving three languages often produced L3 interference when the L2 was stimulated - a phenomenon that is commonly experienced by trilingual people operating in their L2, but not easily explained by the current models.

In this paper, I want to describe our attempts to model some aspects of bilingual aphasia using random boolean networks. Our original intention had been to attempt to construct networks that were direct analogues of some of the cases described in Albert and Obler's book (Albert and Obler 1978). A&O's chapter 4 contains a set of 108 case studies, described in some detail, and I thought that it might be possible to develop network models that mimicked some of the gross features of a few of these cases. Modelling specific cases may seem like an ambitious goal: it is obvious, however, that random networks are an extremely powerful descriptive device, and in principle, it ought to be possible to engineer a network so that it displayed a set of desired characteristics. We have, in fact been exploring this type of problem using a genetic algorithm approach. In practice, however, modelling A&O's cases turned out to be more difficult than we had anticipated. The main reason for this was not that it was difficult to engineer the networks, however. Rather the problem was that the case studies found in A&O were described too loosely for our purposes. In most of the descriptions, for instance, it is hard to tell how long it took for the recovery to occur, and the nature of the remaining deficits is similarly vague. These problems - which are by no means confined to A&O's work – meant that it was difficult to specify what our models were required to do. Another, different approach was clearly necessary.

I mentioned earlier that random bilingual networks allowed to stabilise typically settle down into a state where one of the two languages dominates the other - the dominant language has a greater number of ON items than the non-dominant language does. We can use the number of ON units in each of the two languages as a simple shorthand measure of the current state of the network. Using this metric then suggests an obvious way of exploring the behaviour of bilingual networks under stress: we can "traumatise" a network by making sudden, large changes its the current state description, and we can then explore how the network copes with this. One of the easiest ways of traumatising a network in this way is simply to turn OFF all the words simultaneously. Left to itself, a network in this state will stay there: there is no internal stimulation in the network which would cause it to move to a new state. However, if we "kick-start" the network, by turning on a small number of words at random, this is often sufficient to set the network back on the road to recovery. What we are interested in, in this paper, is whether anything like "patterns of recovery" can be found in the way networks respond to this treatment.

In order to examine this question, I ran a set of some 25 simulations. Each simulation consisted of the following steps. First, a random network was initialised and allowed to stablilise. Each network consisted of two sets of 80 "words". Each word was connected randomly to two other same-set words, with the exception that 10% of the units had one link to both sets of words. These cross language units are Boolean EOR units - they go ON if one or other of their input units is ON, but not if both are ON. All other units were randomly set as AND units - these go ON if both their inputs are ON - or OR units - these units go on if either or both of their input words is ON. Next, the network was "traumatised", by having ALL of its words set to OFF. Finally, four randomly chosen units in one of the two languages were switched ON, and the network was left to stabilise.

We can observe the process of recovery in these networks by counting the number of words which are in the ON state at each cycle of the stabilisation process. This is something of a simplification, of course; there is a lot more going on in these networks than this data implies, but the number of ON items is a useful shorthand at this stage. Some data of this kind has been reproduced in figures 1 to 7.

Figure 1

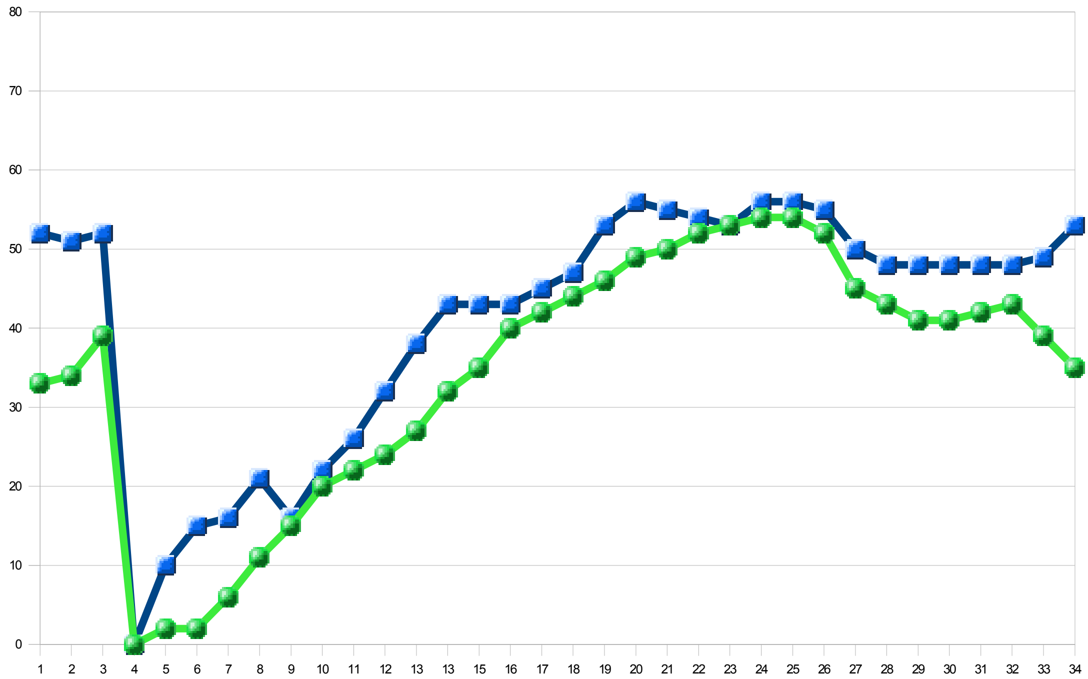

Figure 2

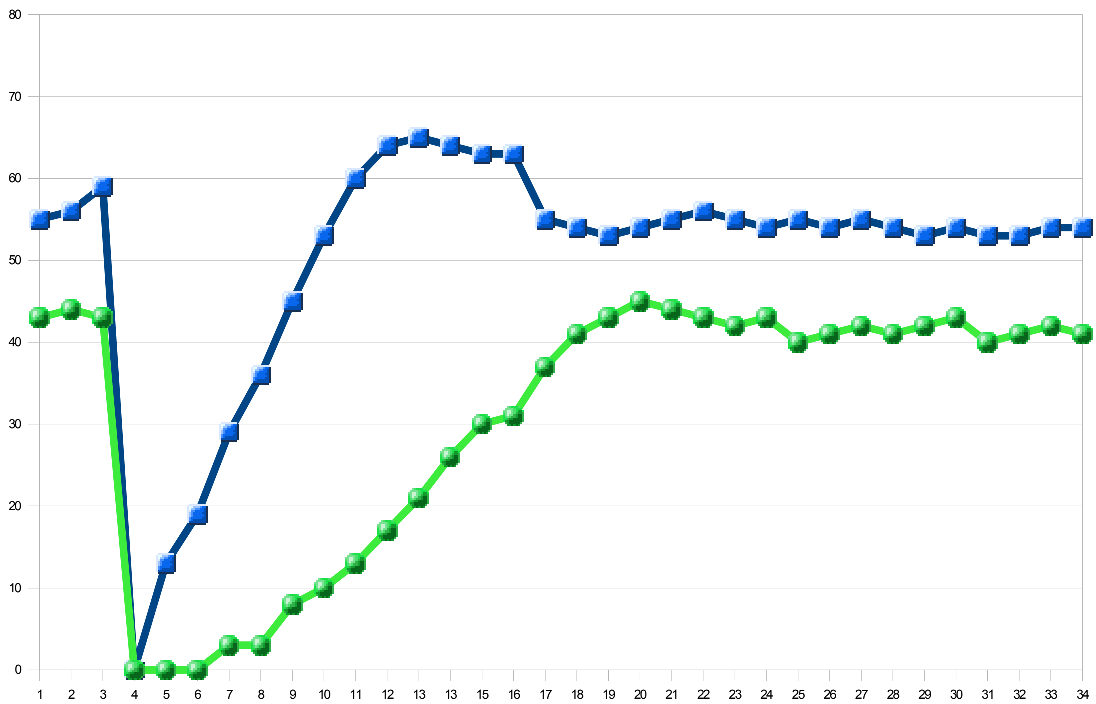

Figure 3

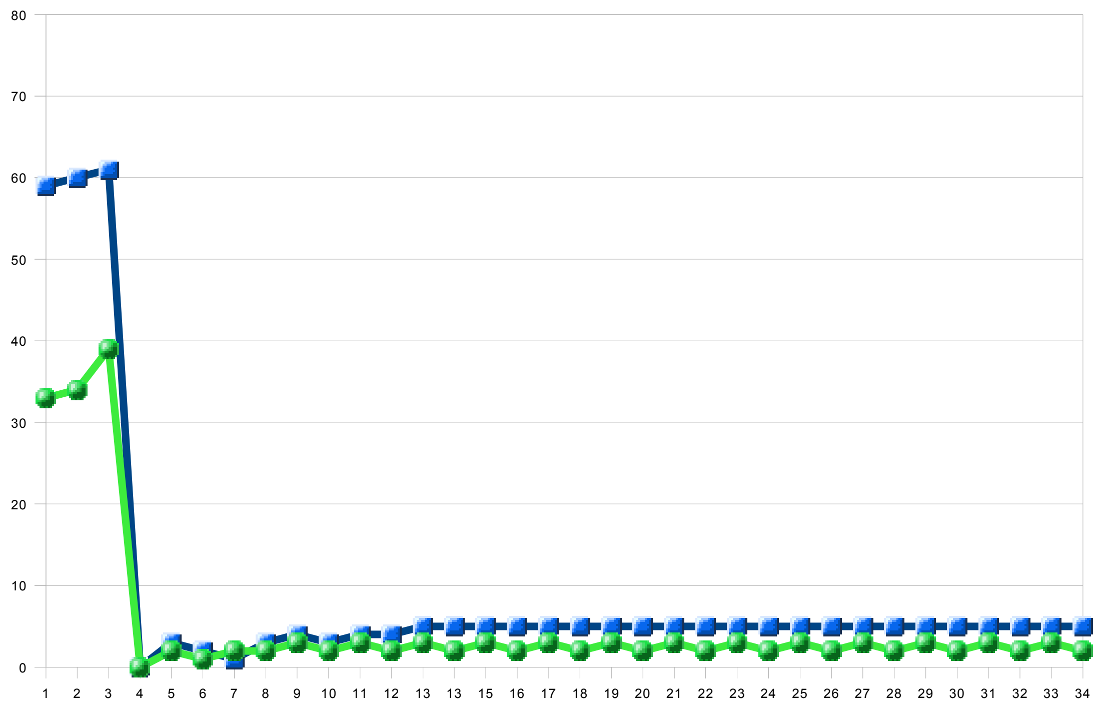

Figure 4

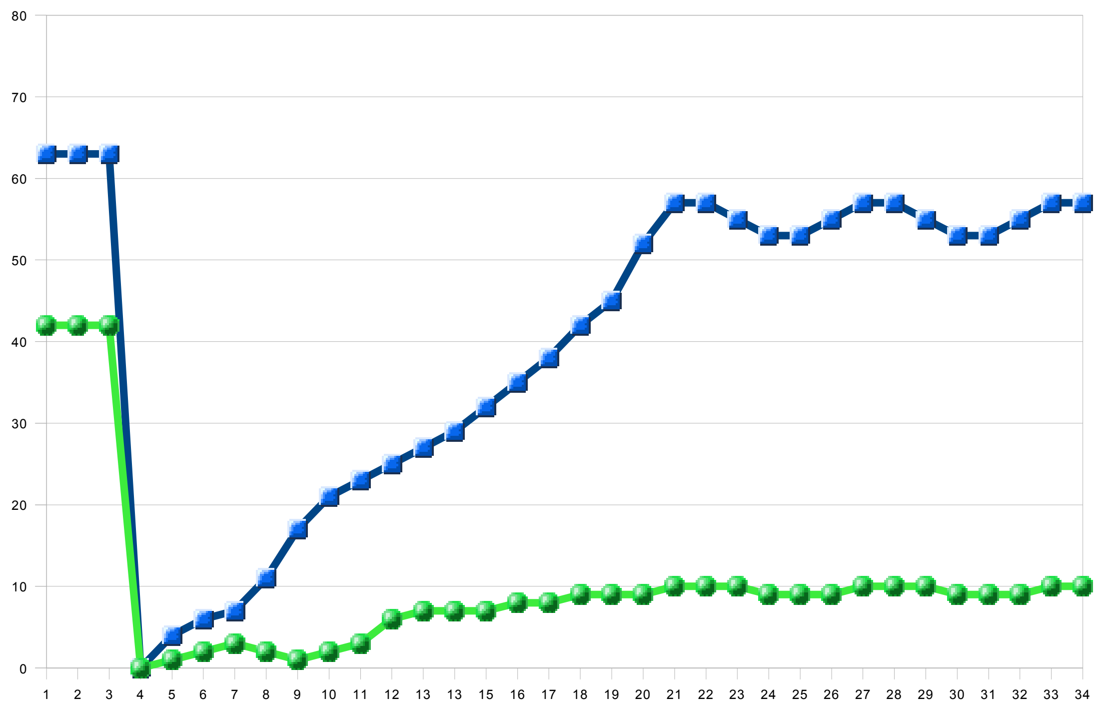

Figure 5

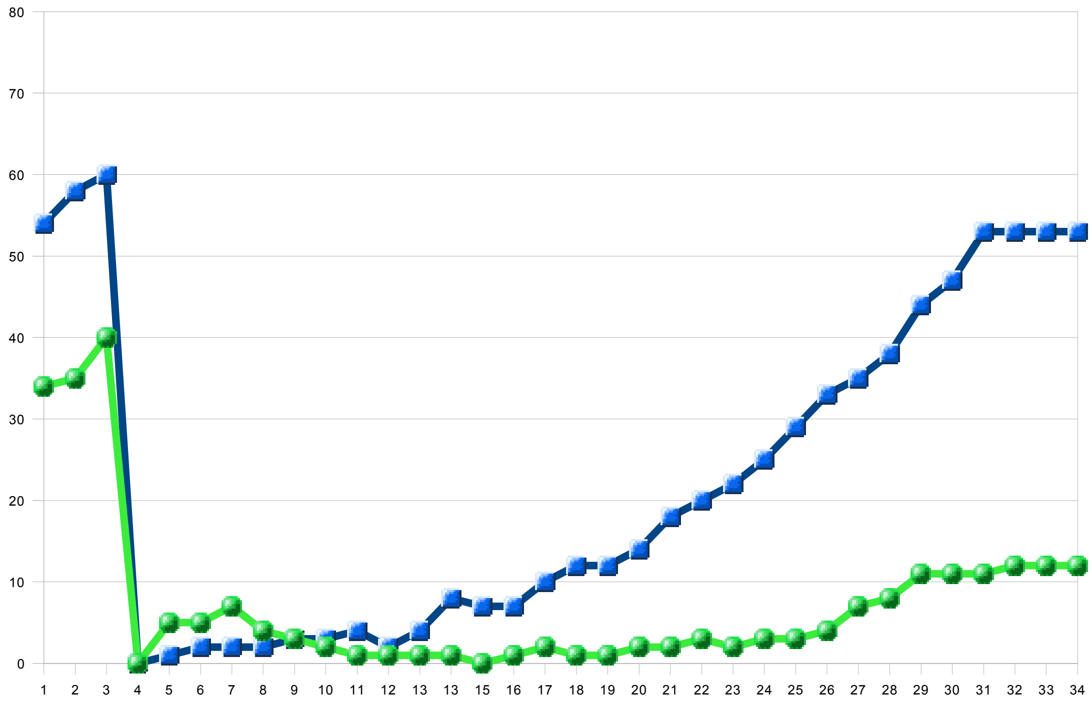

Figure 6

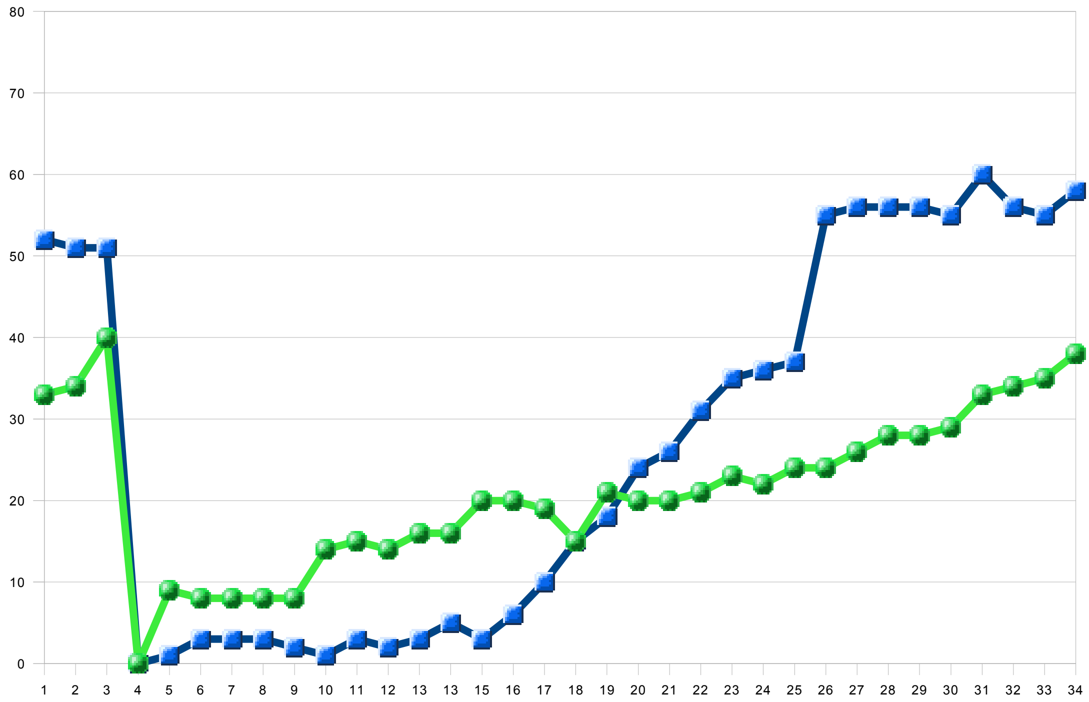

Figure 7

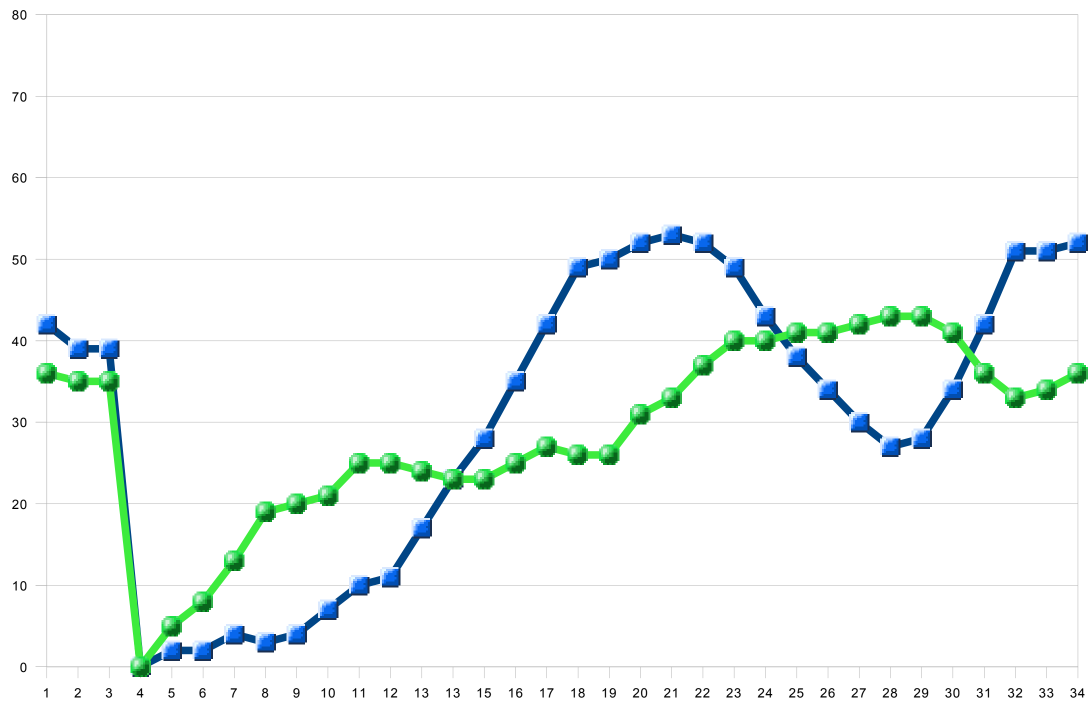

The main thing to note about this data is that several of the recovery patterns appear to be relatively "well-behaved". From the initial kick-start, both languages appear to recover well. In figure 1, for instance, we have a network in which both languages move quickly to their initial equilibrium level. L1 appears to recover slightly more rapidly than L2, and but L2 largely follows the same recovery pattern as L1. In figure 2, both languages also reach their initial equilibrium level, although in this case L2 appears to take rather longer to recover than L1 does. Figure 3 is also "well-behaved" in that both languages fare equally badly. In this case, neither language gets beyond a minimal level of recovery. Figure 4 shows a system which does not return completely to its original equilibrium level. L1 appears to recover quite well, but eventually settles into a pattern of activity which is slightly lower than the original equilibrium level. L2, on the other hand shows only a minimal level of recovery. The remaining figures show somewhat less well-behaved patterns of recovery. Figure 5 shows a case where the first language to show signs of recovery is the L2. This early promise fails to materialise, however: L1 begins to recover slowly, and eventually reaches an equilibrium level not far short of the original level. L2 appears to do nothing until the L1 is almost completely re-established, but it manages to reach a low-level equilibrium after the stabilisation of L1. Figure 6 shows another case where L2 appears to recover before L1. This recovery, though fairly slow, continues until L2 reaches its original equilibrium level. Meanwhile, L1, which appears not to be recovering at all in the early stages of the stimulation, suddenly begins to recover very rapidly, and quickly comes to reassert its dominance again. The most intriguing recovery pattern is figure 7. In this simulation, we seem to have an oscillating pattern of recovery in which the two languages appear to take turns in dominance. L2 appears to be recovering quickly at first, but the rate of recovery slows after the L1 items start to become active. Meanwhile, the recovery pattern for L1 is markedly different: a period of rapid recovery being followed by periods of substantial loss, which allows the L2 to re-establish temporary dominance.

What does all this mean? It would be very easy to suggest that the patterns of recovery shown in figures 1 to 7 are direct analogues of cases reported in the clinical literature. Paradis (1977) for instance, in a large-scale review of case studies, identified five principal patterns of recovery in bilingual aphasia. These were *Parallel Synergistic Recovery*, where both languages recover at roughly the same rate; *Differential Synergistic Recovery*, where both languages recover, but one language recovers faster than the other; *Antagonistic Recovery,* where recovery in one language is achieved at the expense of the other; *Successive Recovery,* where one language recovers fully before any recovery occurs in the other language; and *Selective Recovery*, where a patient fails to regain one of his languages at all.

All five recovery patterns are present in the simulations reported here. Figure 1, for instance, looks like a case of complete and parallel recovery. Figure 2 looks like a case of complete recovery for L1, but delayed recovery for L2. Figure 3 looks like a case of permanent non-recovery. Figure 4 looks like a case of near-total recovery for L1, and severe permanent impairment for L2. Figure 5 looks like a case of very slow recovery in L1. Figure 6 looks like a case of early recovery in L2 followed by delayed but complete recovery in L1. Figure 7 is clearly a case of antagonistic recovery, with L1 and L2 alternating as to which of them is dominant. The parallels between these simulations and some of the cases reported in the clinical literature are striking. However, it would clearly be a mistake to push the parallels too hard. The simulations are NOT intended to be real models of real patients, and the device of using the number of ON units as a short-hand way of describing the state of the systems is clearly a crude simplification.

Nonetheless, the fact that it has been possible for us to produce something that bears more than a passing resemblance to what appears in the clinical literature ought to give us some pause for thought. It is generally assumed that the range of recovery patterns shown by real aphasic patients is a reflection of the enormous complexity of the way language is organised in the human brain, and a great deal of effort has gone into trying to decipher the complexities of this organisation by examining recovery patterns in detail and depth. We have, for instance, *Pitre's Rule*, which states that the most "familiar", or most recently used language returns first, and *Ribot's Rule* which suggests that the first learned language should recover first, and a number of other plausible attempts to explain the variation in recovery patterns found in real aphasic patients. Albert and Obler (1978:106) summarise despairingly, with the suggestion that "no single rule can yet predict language recovery patterns in all the individual cases of prolonged aphasia". The key word here is "yet", implying as it does, that more case studies, and more detailed biodata might be able to show us what this single rule could be. The importance of the simulations I have reported here is that they suggest this hope may be a vain one. What these simulations show is that complex patterns of recovery do not necessarily imply a complex underlying organisation. The recovery patterns reported here were produced by randomly organised networks, with minimal interconnections, and barely enough internal structure to distinguish them from a mere list of words. Nor were these patterns difficult to find. The seven I have reported here were drawn from a batch of 25 simulation runs, with the other runs producing patterns that looked similar to these. This suggests to me that the range of cases reported in some of the clinical literature may in fact be less interesting than it looks at first sight. If complex patterns of recovery can be produced by nothing more structured than a small-scale, random network, then it is perhaps not surprising to discover that real human aphasics, whose lexicons are presumably a lot more complex than these random models, generate a very wide range of recovery patterns too. If this view is correct, then a very large number of different recovery patterns is precisely what we would expect to find in a large group of patients. New cases do not necessarily add anything deep to our understanding of the way languages can recover from trauma. They may merely confirm that language is complex, and produces complex patterns of behaviour.

The question we ought to be asking is what is it about the way language is structured that allows it to recover from severe trauma at all. The simulations reported here suggest that the ability to recover from trauma might be an emergent property of any network structure with minimal organisation. The implications of this conclusion are not at all clear. However, it does suggest that simulations of this type, despite their crude simplifying assumptions, might be able to throw some light on why real patients show such a large range of recovery patterns, and why this data has so far defied any real analysis.

References

Albert, M and L Obler
The bilingual brain. London: Academic Press. 1978.
de Bot, K.
A bilingual production model: Levelt's "speaking" model adapted. *Applied Linguistics* 13(1992), 1-24.
de Groot, AMB.
Determinants of bilingual lexico-semantic organisation. *Computer Assisted Language Learning* 8(1995) , 151-180.
Dijkstra, T and K de Smedt.
Computational Psycholinguistics. Hove: Lawrence Erlbaum Associates. 1996.
Grendel, M.
Verlies and herstel van lexicale kennis. Doctoral Thesis, Nijmegen University. 1993.
Kauffman, SA.
Antichaos and Adaptation. *Scientific American* 265(1991), 64-72.
Kauffman, SA.
The origins of order. Oxford: Oxford University Press. 1993.
Levelt, WJM.
Speaking. Cambridge, Ma.: MIT Press. 1989.

Meara, PM.
Self-organisation in bilingual lexicons. In: **P Broeder and J Murre** (eds.) *Language and Thought in Development: cross-linguistic studies.* Tuebingen: Gunter Narr Verlag. 1996.
Paradis, M.
Bilngualism and Aphasia. In: **H Whitaker and H Whitaker** (eds.) *Studies in Neurolinguistics vol 3.* New York: Academic Press. 1977.
Schreuder, R and B Weltens (eds.)
The bilingual lexicon. Amsterdam: Benjamins. 1993.
Weltens, B and M Grendel.
Attrition of vocabulary knowledge. In: **R Schreuder and B Weltens** (eds). 1993. *The bilingual lexicon.* Amsterdam: Benjamins. 1993.

This paper first appeared in *International Journal of Bilingualism* 3,1(1999), 45-54.

Appendix One: Random Boolean Networks.

This figure shows a very small network consisting of only four units, labelled A,B, C, and D. Each unit can either be ON or OFF. Each unit is connected to two other units, and the current state of each unit is determined by the state of its two input units. In this example,

A is connected to B and C
B is connected to C and D
C is connected to D and A
D is connected to A and B

Three units (A,B and C) are EOR units: that is
A goes ON if either B or C is ON or if both B and C are ON.
B goes ON if either C or D is ON or if both C and D are ON.
C goes ON if either D or A is ON or if both D and A are ON.
The remaining unit (D) is an AND unit: that is
D goes ON only if both A and B are ON.
These values and connections are assigned at random.

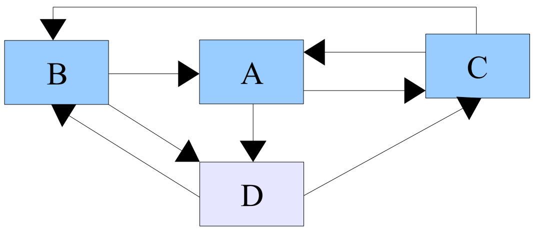

When the network is initialised, units are turned ON or OFF at random. From this initial starting configuration, the pattern of activity in the network can be worked out from a simple rule table. In this case, there are sixteen possible states, and the the following rules apply (1=ON, 0=OFF):

	A	B	C	D	A	B	C	D	
0	0	0	0	0	0	0	0	0	0
1	0	0	0	1	0	1	1	0	6
2	0	0	1	0	1	1	0	0	12
3	0	0	1	1	1	1	1	0	13
4	0	1	0	0	1	0	0	0	8
5	0	1	0	1	1	1	1	0	13
6	0	1	1	0	1	1	0	0	14
7	0	1	1	1	1	1	1	0	14
8	1	0	0	0	0	0	1	0	2
9	1	0	0	1	0	1	1	0	6
10	1	0	1	0	1	1	1	0	14
11	1	0	1	1	1	1	1	0	14
12	1	1	0	0	1	0	1	1	11
13	1	1	0	1	1	1	1	1	15
14	1	1	1	0	1	1	1	1	15
15	1	1	1	1	1	1	1	1	15

What this table shows is that if the system starts off in state number 6 (A=OFF, B=ON, C=ON, D=OFF), then it will move to state 14 (A=ON, B=ON, C=ON, D=OFF). Typically, networks of this sort tend towards a steady state. In this case, the network has two steady states, state 0 (A=OFF,B=OFF,C=OFF,D=OFF) and state 15 (A=ON, B=ON, C=ON, D=ON). Initialising the network in any of the other states will cause it to move to to state 15. For example, initialising the network in state 4 (A=OFF, B=ON, C=OFF, D=OFF) will cause the network to move through the following sequence of states:

	A	B	C	D
state 4	0	1	0	0
state 8	1	0	0	0
state 2	0	0	1	0
state 12	1	1	0	0
state 11	1	0	1	1
state 14	1	1	1	0
state 15	1	1	1	1

State 15 is called an Attractor. In this state the network is stable, and resists attempts to move it to another state.

Larger networks have more interesting properties than a small network of this sort. With a network of any size, it becomes practically impossible to enumerate all the possible states. Typically, however, a randomised network consisting of, say, 200 units, when allowed to stabilise, will reach one of its attractor states in fewer than 30 iterations. Typically, the attractors of larger networks contain a significant number of OFF units.